Ergodicity Economics – Valid Criticism of Traditional Economic Theory or an Imprudent Diatribe?

Luca Formenti

February 24, 2021

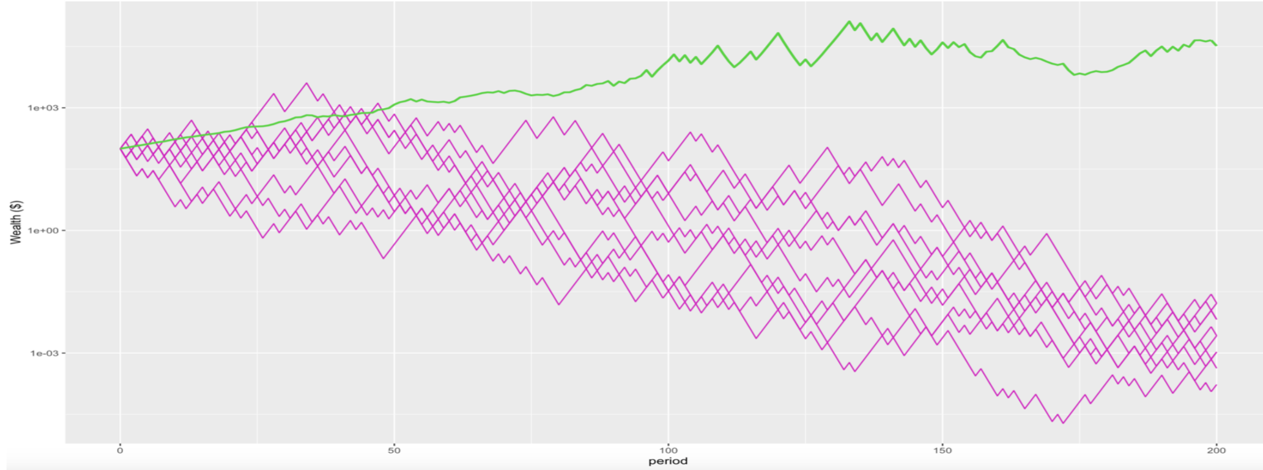

Figure 1: The wealth of ten randomly selected individuals compared to the group average.

If a random process is followed for a sufficiently long time, one is eventually going to see every possible outcome, and the probabilities of these individual outcomes can be computed. This is, in essence, what ergodic theory is all about; a branch of statistics that deals with the probabilistic properties of dynamic systems. Generally, ergodic theory investigates systems that possess the properties of ergodicity. The reasoning behind it is that if a process is observed for a long period of time, one can determine all there is to know about that process in terms of statistical behavior because, in the limit, what one observes converges to the "true" properties of the process. In more technical terms, an ergodic variable satisfies the property that its time average equals, in the limit, its expectation:

$$\lim_{x\to\infty} \frac{1}{T}\int_0^T [f(\omega(t))dt] = \int_{-\infty}^{\infty} [f(\omega)P(\omega)d\omega]$$

For example, if a million people roll a dice simultaneously, add up all the points and divide the result by one million, it would yield the same result (the expected value of 3.5) as if an individual would roll that dice a million times himself, add up the points, and divide the result by one million. But how does this abstract and very technical field of physics relate to the discipline of economics? In essence, the criticism that ergodicity economics brings up is precisely that classical economics always assumes the ergodicity of a system – even in cases where such an assumption models a certain process in an entirely inappropriate way (this assertion has been rejected by many economists; more about this later). In particular, it scrutinizes expected utility theory, which states that people maximize their expected utility – the outcomes of certain events weighted by the respective probabilities – and then behave accordingly. In order to illustrate this and provide more intuition behind the equation above, let us walk through the conventional example to explain non-ergodicity. Say a game is played where the participants flip a coin and start with a wealth of \$100. They lose 40% of their money if the tail comes up, but win 50% if the head appears. They would be required to play this game for a week, flipping a coin every minute. Based on the immediate intuition, is this game worth playing? Proponents of ergodicity economics assert that according to classical economics, one would always play the bet. Since a fair coin is assumed, the expected value of this game is positive, i.e. in the limit, there should be a net increase in the initial wealth. Interestingly, when people are faced with this proposal in real life, they often decline to play. Again, ergodicity economics alleges that a classical economist would say that this is irrational behavior, as the expected value is obviously positive. Yet, under closer inspection, this behavior is not irrational. The alleged problem with the classical economic view is that it holds true only in the aggregate, but not for the individual. Let's get a million people to play this game. When computing the average monetary gain over all players, the absolute wealth does indeed increase, as classical economic theory suggests. However, when looking at the individual, an entirely different picture is seen. Since this game is non-ergodic, the expected value of the group of a million people will differ from the time average for a certain individual. Actually, an individual will lose, in the limit, approximately 5.13% ($1 - [1.5 * 0.6]^{\frac{1}{2}}$) of his or her wealth in each period. The fact that this time average and the group average diverge proves the non-ergodicity of this game. The reason behind this discrepancy between the two averages is that when considering the group average, there will be a few incredibly lucky individuals that would be able to accumulate vast amounts of wealth, while the majority of the group ends up with virtually nothing. Hence, the group average will get distorted by these statistical outliers, while the time average tracks what usually happens to a given individual. I illustrated this divergence through a simulation in R that tracks what happens to the wealth of ten randomly selected individuals compared to the group average: As one can clearly see in Figure 1, the intuitive reaction of many people to not play the game does not seem so irrational anymore. After all, the probability of being one of the lucky few that manage to accumulate millions of dollars (notice the log scaled graph) is ridiculously small. Most likely, a typical person will end up belonging to one of the many individuals that will eventually lose everything. What is important to realize when looking at this picture is that the constraint of 200 periods is just to exemplify the point of non-ergodicity. Of course, when this game is played long enough, the entire wealth of the group will collapse to zero because each individual is most likely to lose 5% on average in each period, as already seen. However, it is still worth taking note of the fact that the intuitive behavior that is perceived as "irrational" by classical economists (at least this is what advocates of ergodicity economics claim) is actually the right thing to do. The takeaway that ergodicity economics proposes for economic modeling is that where appropriate, i.e. when one is faced with non-ergodicity, one should supposedly replace expectations with time averages in order to be able to accurately represent the intertemporal consequences. Whether this is a valid suggestion will be discussed later.

Now, how can this knowledge be utilized when faced with, for example, an investment decision? To understand this, the game needs to be modified a little bit to provide for the possibility of playing the game with only a fraction of the initial wealth. What essentially caused the fast decline in wealth for the majority of the individuals in the game was that they were fully leveraged, i.e. they were invested with their whole wealth ($100). It turns out that by investing only a fraction of the initial wealth into this game, one can actually make it profitable. In portfolio theory, there is a so-called Kelly strategy, which provides for the share of wealth one should invest that almost surely leads to a higher payoff when compared to any other strategy one could take. The Kelly strategy takes a very simple form, namely:

$$f^* = \frac{p(b+a)-1}{b}$$

The share of wealth that should be invested (f^*) is determined by the net odds b, and the probability of winning p. When computing the Kelly strategy for the case at hand ($p = 0.5, b = \frac{0.5}{0.4}$) the result depicts that 10% of the wealth should be invested in order to maximize the growth rate of the initial wealth in the long run. While there is still a large wealth inequality within the group when playing the game with these modified rules (most people still end up below the average), after a sufficiently large number of periods, almost everyone will have increased their wealth. Hence, when keeping the economic framework of optimization in the context of this game, it could be stated that people should be modeled as maximizing growth rates rather than utility itself.

There also has emerged significant criticism in response to ergodicity economics. Doctor, Wakker and Wang (2020)[1] point out that the critique provided by ergodicity economics is largely unjustified, and provides a false picture of the traditional assumptions made in economics. The first and most obvious response to ergodicity economics is that the actual assumption of ergodicity is not implicitly made in expected utility theory. Expected utility theory concerns static decisions under uncertainty, and does not conjecture anything about a time dimension. Further, they argue that ergodicity economics tends to interpret incoherence in empirical experiments on expected utility as a disproof of the theory – even though the theory has been constantly improved precisely on the basis of such inconsistencies in the past. In essence, the case presented in ergodicity economics is merely an exceptional instance that comes down to the Kelly criterion outlined above. In this sense, they suggest that the proponents of ergodicity economics should first dig deep into the literature of basic economics before formulating an injudicious denunciation of it.

To conclude, ergodicity economics seems to provide interesting insights and inspires reflection. Yet,

[1]Doctor, J., Wakker, P., & Wang, T. (2020). Economists' views on the ergodicity problem. Nature Physics, 16(12), 1168-1168. doi: 10.1038/s41567-020-01106-x

under closer scrutiny, the theory seems indeed to be a fallacious attack of traditional economic theory, based on erroneous assumptions and a lack of understanding of the relevant literature that is already out there.